



**Europäisches
Patentamt**

**Eur pean
Patent Office**

**Office européen
des brevets**

Bescheinigung

Certificate

Attestation

Die angehefteten Unterla-
gen stimmen mit der
ursprünglich eingereichten
Fassung der auf dem näch-
sten Blatt bezeichneten
europäischen Patentanmel-
dung überein.

The attached documents
are exact copies of the
European patent application
described on the following
page, as originally filed.

Les documents fixés à
cette attestation sont
conformes à la version
initialement déposée de
la demande de brevet
européen spécifiée à la
page suivante.

Patentanmeldung Nr. Patent application No. Demande de brevet n°

02102550.7

Der Präsident des Europäischen Patentamts;
Im Auftrag

For the President of the European Patent Office

Le Président de l'Office européen des brevets
p.o.

R C van Dijk



Anmeldung Nr:
Application no.: 02102550.7
Demande no:

Anmeldetag:
Date of filing: 08.11.02
Date de dépôt:

Anmelder/Applicant(s)/Demandeur(s):
International Business Machines Corporation

New Orchard Road
Armonk, NY 10504
United States of America
US United States of America

Bezeichnung der Erfindung/Title of the invention/Titre de l'invention:
(Falls die Bezeichnung der Erfindung nicht angegeben ist, siehe Beschreibung.
If no title is shown please refer to the description.
Si aucun titre n'est indiqué se référer à la description.)

Method And System For Reducing Host Variable Impact On Access Path Selection

In Anspruch genommene Priorität(en) / Priority(ies) claimed / Priorité(s)
revendiquée(s)
Staat/Tag/Aktenzeichen/State/Date/File no./Pays/Date/Numéro de dépôt:

Internationale Patentklassifikation/International Patent Classification/
Classification internationale des brevets:

G06F17/30

Am Anmeldetag benannte Vertragsstaaten/Contracting states designated at date of
filing/Etats contractants désignées lors du dépôt:

AT BE BG CH CY CZ DE DK EE ES FI FR GB GR IE IT LI LU MC NL PT SE SK TR

D E S C R I P T I O N

Method And System For Reducing Host Variable Impact On Access
Path SelectionBackground of the Invention

1. Field of the Invention

This invention relates in general to computer-implemented database systems, and, in particular, to improving an access path selection for Structured Query Language with variables in a relational database management system.

2. Description of the Related Art

Databases are computerized information storage and retrieval systems. A Relational Database Management System (RDBMS) is a database management system (DBMS) which uses relational techniques for storing and retrieving data. Relational databases are organized into tables, which consist of rows and columns of data. The rows are formally called tuples. A database will typically have many tables and each table will typically have multiple tuples and multiple columns. The tables are typically stored on random access storage devices (RASD) such as magnetic or optical disk drives for semi-permanent storage.

RDBMS software using a Structured Query Language (SQL) interface is well known in the art. The SQL interface has evolved into a standard language for RDBMS software and has been adopted as such by both the American National Standards Institute (ANSI) and the International Standards Organization (ISO). The SQL interface allows users to formulate relational operations on the tables either interactively, in batch files, or embedded in host languages, such as C and COBOL. SQL allows the user to manipulate the data.

The definitions for SQL provide that a RDBMS should respond to a particular query with a particular set of data given a specified database content, but the method that the RDBMS uses to actually find the required information in the tables on the disk drives is left up to the RDBMS. Typically, there will be more than one method that can be used by the RDBMS to access the required data. The RDBMS will optimize the method used to find the data requested in a query in order to minimize the computer time used and, therefore, the cost of performing the query.

The RDBMS software determines an access path for an SQL statement during a bind process. SQL compilation is called a BIND process, and the output of the BIND process is a plan, which is a compiled run-time structure used for executing the SQL statement. The plan includes access paths, which are the paths the RDBMS uses to get to the data that SQL statements request. During the BIND process, an optimizer of the RDBMS software selects an access path for a SQL statement. The access path is the key to determining how well an SQL statement performs.

Most RDBMS such as DB2, Informix, Oracle and MS SQL Server implement so called cost-based optimizer. The optimizer is a component that, based on the application request specified in form of SQL statement, selects, i.e., generates optimal access path to the requested data. The access path details are stored in a form of an executable code either on disk (for static SQL) or in the processors cache (for dynamic SQL). This code is referenced during subsequent statement executions. This way the RDBMS saves significant processor resources because, in general, the optimization process is many times (for a typical Online Transaction Processing (OLTP) workload, 100s of times) more expensive than the statements execution itself.

In order to select the optimal access path, the cost-based optimizers (unlike rule-based ones) take into account the current status of some environmental variables. The most important of these variables is the statistics such as cardinality, size and frequency distribution of the database objects that need to be accessed. While the cardinality (the number of distinct values a database object has) and its size can always be available to the optimizer, the frequency distributions of table columns' values are often unknown. The reason is that the so-called host variables regularly change with each statements execution.

US 6,006,220 by Donald J. Haderle et al., assigned to IBM Corporation, Armonk, N.Y., US, filed Sep. 30, 1997, issued Dec. 21, 1999, "Determining the optimal access path for a query at execution-time using an actual value for each variable in a query for estimating a filter factor" states that, currently, when an SQL statement includes a variable, the default filter factor, also known as selectivity, selected for the SQL statement is typically not optimal, and, therefore, the access path that is determined is not necessarily optimal. Therefore, a method is suggested of executing a query in a computer, the query being performed by the computer to retrieve data from a database stored on a data storage device connected to the computer, the method comprising the step of: determining an optimal access path for the query at execution time using an actual value for each variable in the query to estimate a filter factor, wherein each variable is a special register.

Object of the Invention

Starting from this, the object of the present invention is to provide a method and a system that facilitates a more efficient execution of database queries containing variables.

Brief Summary of the Invention

The foregoing object is achieved by a method and a system as laid out in the independent claims. Further advantageous embodiments of the present invention are described in the sub claims and are taught in the following description.

According to the present invention a method and a system is provided for processing a query in a database management system. Firstly, a preferred access path is generated for a query at execution time on the basis of a first value for each variable in the query. The information about said query, about said first value for each variable and about said preferred access path is stored. In case the same query is to be executed again using a second value for each variable, said information about said first value for each variable is compared with information about said second value for the respective variable and it is determined whether or not a re-generation of the preferred access path is to be performed again, and, in the affirmative, a re-generation of the preferred access path is initiated.

The step of re-generating the preferred access path may include the step of storing the newly generated preferred access path, the associated host variables and frequency distributions. In other words, when a re-optimization is needed, the new access path would not overwrite the old one, but just be added to a list of possible access paths.

Hence, the present invention proposes enhancing the access path selection process by storing the relevant frequency distribution and the host variables values used for selecting the access path together with the corresponding access path executable. This way the host variables impact to the access path can be

inexpensively assessed and the re-optimization triggered only if necessary.

Let's demonstrate the idea in an example.

Let's T(C1,C2, C3) be a table with columns C1, C2 and C3, and an index (C1,C2).

card(T) = 1000.

card(C1) = 3. The values are 'A', 'B' and 'C'.

card('A') = 500, card('B')=490, card('C')=10

The following code segment needs to be executed:

```
hv = value for C1 specified as a host variable (e.g. user
input via some data entry facility)
SELECT * FROM T WHERE C1 = hv
```

Assume that the values assigned to hv are 'B', 'A', 'D'. The access path selection process, i.e., the optimization, is delayed to statement execution time. For the first statement execution the optimizer will generate a tablespace scan access path (because cardinality of 'B' is very high) and store it for subsequent executions. Along with this code the frequency distribution for C1 and value 'B' will be stored. When the statement is executed for the second time, a quick check will reveal that the value of 'A' does not change the access path (because the cardinality of 'A' is close to the cardinality of 'B'), therefore the existing executable can be used without re-optimization. At third execution the 'hv' value is 'D' and the check will signal that the re-optimization needs to be done (because cardinality of 'D' is 0, i.e., very different from its predecessors).

This methodology ensures that the statements are always executed with the optimal access characteristics, but without costly re-optimizations that are currently associated with every statement execution.

The most efficient implementation of the idea includes storing all the different access paths and their associated host variables values and frequency distributions as they get generated. In other words, when a re-optimization is needed, the new access path would not overwrite the old one, but just added to the list of possible access paths.

Brief Description of the Several Views of the Drawings

The above, as well as additional objectives, features and advantages of the present invention, will be apparent in the following detailed written description.

The novel features of the invention are set forth in the appended claims. The invention itself, however, as well as a preferred mode of use, further objectives, and advantages thereof, will best be understood by reference to the following detailed description of an illustrative embodiment when read in conjunction with the accompanying drawings, wherein:

Fig. 1 shows a flowchart of a first embodiment of the method in accordance with the present invention;

Fig. 2 shows a detailed flowchart of a first embodiment of the method in accordance with the present invention; and

Fig. 3 shows a block diagram illustrating a database management system in a computer system in accordance with the present invention.

Detailed Description of the Invention

With reference now to Fig. 1, there is depicted a flowchart of a first embodiment of the method in accordance with the present invention, in particular, Fig. 1 illustrates a method of processing a query in a database management system (block 100).

When a query is to be executed a preferred access path is generated for said query at execution time on the basis of a first value for each variable in the query. The first value for each variable is, in this moment, the actual value. Therefore, an optimizer performing the task of generating the preferred access path is able to take into account the actual values of the query variables (block 102).

According to the present invention, information about said query, about said first value for each variable and about said preferred access path is stored for future use (block 104). A unique identifier unambiguously denoting the query to be executed may form the information about the query. Such a unique identifier may be formed by a representation of the query itself, such as the SQL code, or may be generated by some digest function, such as MD5, so that the information about the query may be used as a key in order to retrieve the corresponding information about the value for each variable and about said preferred access path. The information about the value of each variable may be one of the group of, the value itself, a range of values, a category categorizing a range of values, a category categorizing the behavior of the optimizer, i.e., a category of values causing a table-space scan access path, a category of values causing an index-scan access path, a category of values causing an index-access access path, category corresponding to a cardinality range of said first value belongs to or the frequency distribution of the values.

In the following a plurality of different queries may be executed in accordance to the present invention, in particular, according to the steps illustrated by block 102 and block 104 (dotted line 106). However, if the query from the beginning is to be executed again, the method is continued (block 108). It should be noted that a query is considered to be the same query even if the values for the query variables may be different.

Now, the information previously stored and the respective information derived from the current value of each variable of the query is compared (block 110). The comparison may include the generation of the information if they need to be derived from the current variable values. Subsequently, it is determined whether or not the access path needs to be re-generated (block 114). A re-generation of the access path may be necessary, if the current, or second, value of at least one of the variables differs from the stored, or first, value of the same variable by at least a predetermined amount. Alternatively, a re-generation of the access path may be performed, if the category the first value belongs to differ from the category the second value belongs to. In other words, whenever at least one value may cause the optimizer to generate a different access path than the one stored, a re-generation is initiated. It is acknowledged that it may be implemented that a re-generation is only be initiated if more than one of the variable values have been changed.

If yes, the preferred access path is re-generated (block 116) using the current variable values before executing the query according to the newly generated access plan (block 118). In contrary, if no, i.e., it is not necessary to re-generate the access path, the stored access path to be used during execution of the query (block 118) is retrieved (block 120).

With reference now to Fig. 2, there is depicted a detailed flowchart of a first embodiment of the method in accordance with the present invention. After the start (block 200) a new query to be executed is awaited. In a first step a query is received (block 202), then, a unique identification (ID) is derived from as aforementioned with reference to Fig. 1 (block 204).

The query includes SQL statements and variables, also referred to as "host variables". Host variables are variables referenced by embedded SQL statements. They transmit data between the database manager and an application program (cf. Fig. 3). According to the present invention the actual values for the variables are taken into account when optimizing the query for execution. Furthermore, the question whether or not to re-optimize a query, i.e., re-generating a preferred access path, depends on the values used in the creation of an existing access path and the values in the query to be executed. In order to allow comparing different values of each variable in a query, information are generated representing the value, as explained above with reference to Fig. 1 (block 206).

With the unique ID as a key it is determined whether or not information about a query, about a preferred access path and about the value of at least one of the variables is already stored (block 208). In case an access path exists for the looked-up query ID, then, information about the stored query values are retrieved (block 212).

Then, the retrieved information previously and the corresponding information derived from the current value of each variable of the query are compared (block 214). The comparison may be performed as described with reference to Fig. 2 above.

In case a re-generation of the access path is necessary (block 216), it is checked whether or not another access path is stored

under the same query ID (block 218). If yes, the method continues processing at block 212 in order to retrieve information about the query values corresponding to the other access path.

Back to block 216, alternatively, if no, i.e., if a re-generation of the access path is not necessary, be it, because the values are the same, the stored values and the current values belong to the same category or, simply, the current values would not lead to a different access path, the stored access path gets retrieved (block 220), before the query gets executed according to the access path (block 222).

Returning the focus again to block 218, alternatively, if no, i.e., if there is no other access path stored under the current query ID, then, an access path is generated taking into consideration the actual query values (block 224). Subsequently, information about the actual query values and about the generated access path is stored under the query ID for future use (block 226). Then, the query gets executed according to the access path (block 222). Thereafter, the method is continued with the initial step of receiving a query to be executed (block 202).

Now with reference to Fig. 3, there is depicted a block diagram illustrating a database management system 310 in a computer system environment 312 in accordance with the present invention.

The computer system environment 312 further includes an external application 314 that is separate from the database management system 310, a repository 316 and a storage unit 318, both for storing data. An application server, such as SAP, PeopleSoft or Siebel, may form the external application 314. The repository 316 and the storage unit 318 may be formed by non-volatile memory, such as a hard disk. It is acknowledged that the

external application may be running on a different computer system that would be connected to the computer system 312 via a network. The same may apply to the repository 316 and the storage unit 318, which may be formed by one or more dedicated storage servers. Alternatively, the repository 316 could be an integral part of the storage unit 318 or vice versa.

The database management system 310 comprises an optimization unit 322 and other components 324. The other components are illustrated by one single block purely for the sake of clarity. For more details on the operation of a database management system and its components it is referred to Fig. 1 and the respective description.

The optimization unit 322 acts as a device for processing a query in the database management system 310 that generates an access path for retrieving desired data from the computer storage 318 in response to the query. The optimization unit 322 comprises a first interface 332 for receiving a query in form of a predetermined representation, means 334 for deriving an identification from at least a portion of the representation, a second interface 336 for accessing a repository of information about the value of each variable in the query in order to find previously stored information about query values being used to create a stored access path, and a functional unit 338 being adapted to compare the information about said first value for each variable with information about said second value for the respective variable and to determine whether or not a re-generation of the preferred access path is to be performed again, when the same query is to be executed again using a second value for each variable, and if yes, for re-generating of the preferred access path.

The first interface 332 may be formed by a communication link to at least one of the other components 324, whereas the second

interface 336 may be implemented by a communication link to the repository 316.

The means 334 for deriving a unique identification from at least a portion of the representation may be formed by a functional unit, such as a computer program, for computing a hash function, such as MD5, from the respective portion of the query's representation.

In a particular embodiment, the optimization unit 322 enhances a system for the storage and retrieval of data comprising a repository for storing access path information specifying a preferred access path to be used for a particular query, whereby said repository is adapted for storing information of at least one host variable value that has been used to generate said access path information and it contains a functional entity for dynamically optimizing the database access by invoking a re-generation of the access path whenever at least one of said host variable values kept in said repository differs from the actual value by at least a predetermined value, whereby a re-generation of the preferred access path is only performed when the probability of a significant performance loss is higher than the time needed for the re-generation of the predefined information.

The present invention can be realized in hardware, software, or a combination of hardware and software. Any kind of computer system - or other apparatus adapted for carrying out the methods described herein - is suited. A typical combination of hardware and software could be a general-purpose computer system with a computer program that, when being loaded and executed, controls the computer system such that it carries out the methods described herein. The present invention can also be embedded in a computer program product, which comprises all the features enabling the implementation of the methods described herein, and

which - when loaded in a computer system - is able to carry out these methods.

Computer program means or computer program in the present context mean any expression, in any language, code or notation, of a set of instructions intended to cause a system having an information processing capability to perform a particular function either directly or after either or both of the following a) conversion to another language, code or notation; b) reproduction in a different material form.

C L A I M S

1. A method of processing a query in a database management system, the method comprising the steps of:

generating a preferred access path for a query at execution time on the basis of a first value for each variable in the query,

storing information about said query, about said first value for each variable and about said preferred access path,

when the same query is to be executed again using a second value for each variable,

- comparing said information about said first value for each variable with information about said second value for the respective variable and

- determining whether or not a re-generation of the preferred access path is to be performed again, and

- if yes, re-generating the preferred access path.

2. The method according to claim 1, wherein the step of generating a preferred access path includes the step of:

taking into account the frequency distribution of the first value for each variable in the query.

3. The method according to claim 1 or 2, wherein the step of storing information about said query includes the step of:

generating a unique identifier unambiguously denoting the query.

4. The method according to one of the preceding claims, wherein the step of storing information about said first value for each variable includes the step of:

storing said first value for each variable itself.

5. The method according to one of the preceding claims, wherein the step of storing information about said first value for each variable includes the step of:

determining the frequency distribution of said first value for each variable.

6. The method according to one of the preceding claims, wherein the step of storing information about said first value for each variable includes the step of:

determining a unique indicator unambiguously denoting a category of values, the values of which cause the generation of the same access path.

7. The method according to claim 6, wherein the step of determining a unique indicator unambiguously denoting a category of values includes the step of:

choosing one category from the group of, a category of values causing a table-space scan access path, a category of values causing an index-scan access path, a category of values causing an index-access access path.

8. The method according to claim 6, wherein said the step of determining a unique indicator unambiguously denoting a category of values includes the step of:

choosing one category corresponding to a cardinality range of said first value belongs to.

9. The method according to one of the claims 7 or 8, wherein the step of comparing said information about said first value for each variable with information about said second value for the respective variable includes the step of:

categorizing said second value for each variable.

10. The method according to claim 9, wherein the step of determining whether or not a re-generation of the preferred access path is to be performed again includes the step of:

determining whether or not said first and said second value belong to a different category.

11. The method according to one of the preceding claims, wherein the step of re-generating the preferred access path includes the step of:

storing the newly generated preferred access path, the associated host variables and frequency distributions.

12. A system for the storage and retrieval of data comprising:

an optimizer for generating a preferred access path for a query at execution time on the basis of a first value for each variable in the query,

a storage unit for storing information about said query, about said first value for each variable and about said preferred access path,

a functional unit being adapted to compare said information about said first value for each variable with information about said second value for the respective variable and to determine whether or not a re-generation of the preferred access path is to be performed again, when the same query is to be executed again using a second value for each variable, and if yes, for re-generating of the preferred access path.

13. A computer program product stored on a computer usable medium, comprising computer readable program means for causing a computer to perform a method according to anyone of the preceding claims 1 to 11.

A B S T R A C T

Method And System For Reducing Host Variable Impact On Access
Path Selection

The present invention proposes enhancing the access path selection process by storing the relevant frequency distribution and the host variables values used for selecting the access path together with the corresponding access path executable. This way the host variables impact to the access path can be inexpensively assessed and the re-optimization triggered only if necessary. (Fig. 1)

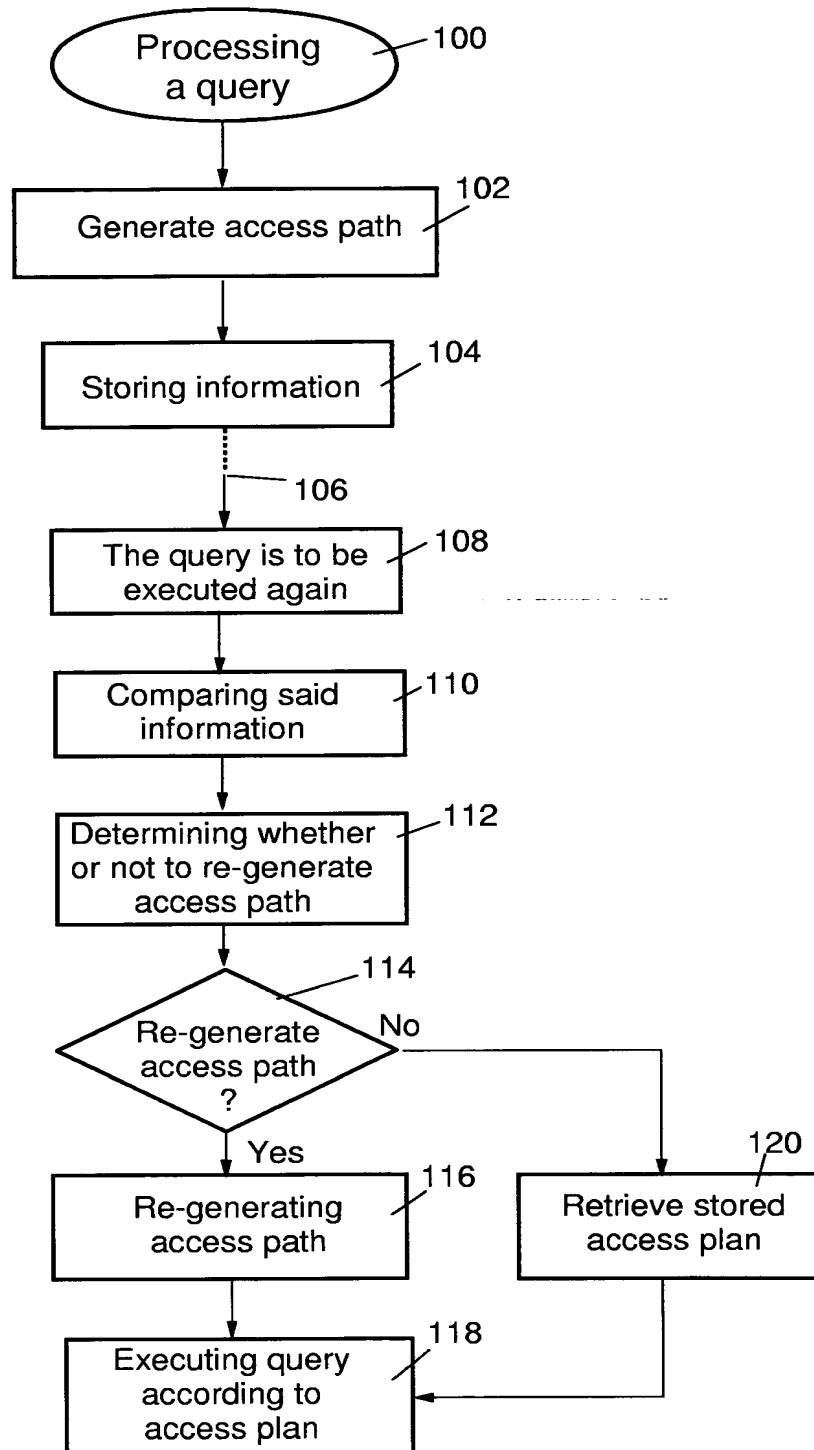


FIG. 1

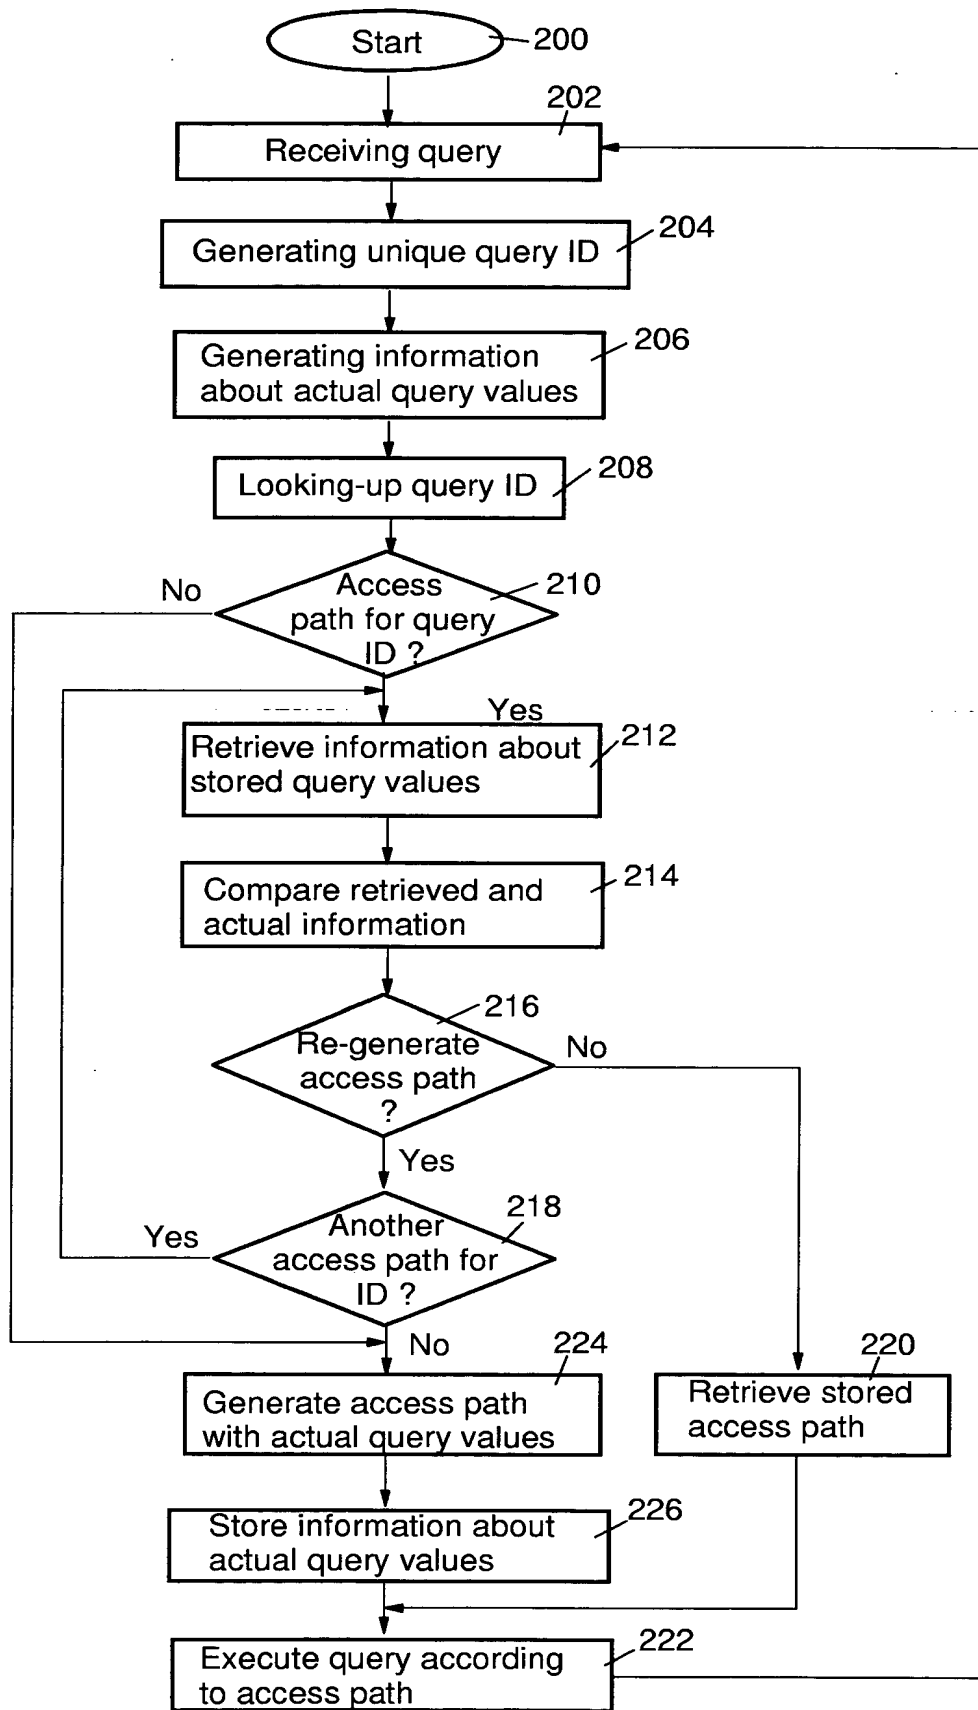


FIG. 2




FIG. 3

